

ABSTRACT OF THE DISCLOSURE

In a technique for compressing data, a data storage server receives a set of data, which includes a first number of subsets. The storage server defines a number of compression groups for the set of data, each compression group including multiple entries, each entry containing a pointer to a different subset of the set of data. At least part of the set of data is then compressed into a smaller number of subsets. For each of the subsets which does not contain compressed data after the compression, a predetermined value is stored in the corresponding entry of the corresponding compression group, to indicate that corresponding data is compressed and represented elsewhere in the compression group. The predetermined value may also indicate which compression algorithm was used to compress the data.